Switched RC Multi-pole Filter

Robert L. Shuler
NASA Johnson Space Center
Avionic Systems Division / EV131
Houston, TX 77058
rshuler@ems.jsc.nasa.gov, 281-483-5258

Abstract – **The design and experimental verification of a switched RC multi-pole filter is presented. This highly compact circuit easily obtains sub-Hz, adjustable response utilizing reasonable sized on-chip components, and multiplexing the main resistor and op amp among filter stages. Design considerations for anti-aliasing, noise avoidance, and dynamic op amp compensation are presented.**

(Note: this circuit concept appears in the literature as far back as 1968, although not recently. The above paper describes a straightforward but potentially useful application of the concept.)

Switched RC Multi-pole Filter

Robert L. Shuler
NASA Johnson Space Center
Avionic Systems Division / EV131
Houston, TX 77058
rshuler@ems.jsc.nasa.gov, 281-483-5258

Abstract - **The design of a switched RC multi-pole low pass filter is presented. This highly compact circuit easily obtains sub-Hz, adjustable response utilizing reasonable sized on-chip components, and multiplexing the main resistor and op amp among filter stages. Design considerations for anti-aliasing, noise avoidance, and dynamic op amp compensation are presented.**

1 Background

Descriptions of switched resistor circuits appear as far back as 1966, shortly after the invention of the integrated circuit, when Lin et. al. [1] describe the application to "sub-audio" low pass filters in IC's. IC's provided the need, due to the area required to fabricate large high quality integrated resistors, and also the opportunity, by increasing speeds and reducing circuit parasitic elements to the degree that such switching might become practical. In 1968 Sun and Frisch [2] provide an analysis for the general case of switched RC networks and use the term "resistor multiplication."

Following in 1969 Kaehler [3] gives an analysis emphasizing the consequent variability of filter parameters as a function of the switching duty cycle, pointing out the flexibility available to share active components (an op amp, in this case), and using the term "capacitor multiplication." Kaehler seems to have control systems in mind, with a band-pass rather than low-pass application, and emphasizes an analysis of phase delay error.

In 1973 Bruton and Pederson [4] change the terminology yet again to "switched conductances" and describe an application to tunable filters for use in seismic analysis or filtering of noisy communication channels. They provide an impulse response analysis. It is interesting to note that none of these authors provides a sampled data analysis. Presumably most of the applications were taken over either by digital filters, or by switched capacitor filters, which began appearing also about this time.

Perhaps the most complete analysis, which does use sample data theory, was a Master's thesis by Paul Embree in 1981-82 [5]. But by May of 1982 a paper entitled "Switched-Resistor Filters..." [6] actually refers to the use of switched capacitors as a resistor, and no further papers regarding true switched resistors appear.

The present author has been seeking monolithic high quality very-low cutoff frequency filters for several years. Some work pushing the state of the art in switched capacitor filters was sponsored

at Texas A&M in 1995 [7], and techniques based on the old idea of using an op amp as a capacitor multiplier were explored [8]. In the process of applying auto-zero offset compensation techniques to low frequency analog circuits, the switched resistor effect was accidentally demonstrated. Subsequent research brought forth the literature described, and a test chip was fabricated to explore the practical usefulness of the idea again with current CMOS circuit technology.

The extremely high switching speeds now available improve the technique by allowing higher sampling frequencies. A technique is described in the current paper for further optimizing the switching ability of a CMOS operational amplifier, and the advantages of sharing the large passive resistor needed for a low pass filter are presented, along with the evaluation of the test chip data. A principal goal will be development of design guidelines for practical low pass filters.

2 Principle of Operation

Consider a network in which a switch is placed in series with a resistor, and a capacitor is available to hold the voltage at a network node for brief periods, such as the low pass filter circuit Figure 1.
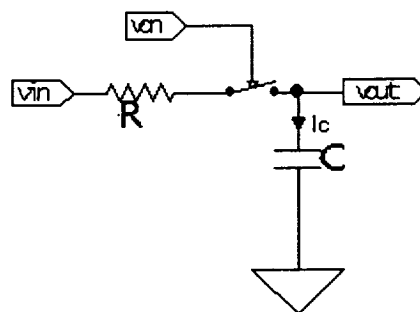

Fig. 1: Idealized switched resistor low-pass filter

Vout will be determined by the integral of the current (Ic) through C. With the switch closed, assuming an ideal switch, and no load at Vout, the current through C will be Ic = (Vin-Vout)/R. With the switch open, of course Ic=0. It is an appealing notion, then, that by operating the switch at a frequency higher than the bandwidth of the signal of interest (Vin), the average current Ic would be reduced by the duty cycle of the switching signal. Thus the time constant of the circuit would be decreased by that amount. Or if we think of the resistor R as defining the current (Ic=(Vin-Vout)/R) then the apparent resistance is increased by the inverse of the duty cycle.

The question of interest then becomes over what operating ranges these approximations yield useful results, and whether actual components provide acceptable approximations to the idea. Each of the references [1] through [5] takes a particular approach to this, usually for more generalized networks.

Since the intended application in this paper is sampled data acquisition, it is observed that by assuming Vin to be sampled in sync with the switching pulses, and that the sampling (and

switching) interval has been chosen by Nyquist criteria to provide acceptable error levels, then a somewhat simpler analysis is possible. This is taken as our first design guideline. It effectively limits the maximum pulse period to be the inverse of the lowest permissible sampling rate.

A second guideline is required which is unique to the switched resistor concept. The "on" portion of the duty cycle will obviously have to be much less than the natural RC time constant. Otherwise the capacitor would simply charge up such that Vout=Vin for each sample, with no practical effect from switching. If duty cycles up to 100% are to be allowed, then this also places a lower limit on sampling rate.

Based on the above two criteria, it seems prudent to make the sampling rate as high as possible. On the high end, the maximum speed of the switches and parasitic circuit effects limit the operation. In the multi-pole circuit described later, an operational amplifier must be switched. The far slower speed of this device limits the maximum rate. The amp must settle and remain accurate for a percentage of the "on" cycle corresponding to the system accuracy desired. Any deviation constitutes an injection of error signal. A technique will be described for improving the speed of the op amp switching.

3 Experimental Verification of Single-Pole Filter

A circuit similar to Figure 1 was fabricated using HP 0.5μ CMOS mixed signal technology. With a silicide block layer available, a value of R = 2 megohms was chosen as practical. The size of the resistor is further improved because precision is not required. If a precise cutoff frequency is needed, it can be trimmed by adjusting the duty cycle. A nominal value of C = 320pF was fabricated using capacitor wells, and Vout was attached to an op amp in unity gain configuration as a buffer to avoid parasitic loading.

Measurements on the test chip indicate both R and C values were slightly larger than planned with a natural cutoff frequency of approximately 200 Hz instead of the planned 248 Hz. Table 1 shows the filter cutoff frequency obtained for various duty cycles using a switching rate of 1000 Hz. The correct single pole response characteristic was verified for input signals up to the switching, i.e. sampling rate, beyond which response is not meaningful. Distortion for frequencies below 20 Hz was basically so low as to not be reliably measureable with our instruments, and total noise and distortion at those frequencies were consistent with what would be required for an overall system accuracy of 12 bits.

Table 1: Filter response vs. duty cycle at 1000 Hz switch rate

cutoff freq. (Hz)	"on" time (μs)	experimental duty cycle (%)	predicted duty cycle (%)
200	1000	100	100
20	100	10	10
5	20	2	2.5
3.3	10	1	1.65
2	5	0.5	1
0.2	0.5	0.05	0.1

The table stops at 0.2 Hz because 0.5μs was the shortest pulse width that could be obtained from our function generator when operating at 1000 Hz. However, simulations indicate that it probably is not reasonable to go much below the 0.5μs "on" time, as the op amp settling time would begin to introduce error components.

There was no obvious reason for the gradual deviation from predicted duty cycles, except that it possibly is within the range of experimental error since the 3db point for filter cutoff frequency was being read from a scope face.

To give a feel for the operation of the filter, we present three scope traces at a reduced switch rate (it is hard to see both switching and input waveform with 2 orders of magnitude difference in time scale) in Figure 3. All traces show response to an input square wave. Traces (a) and (b) also show the switching saveform (low is "on") for different duty cycles and time scales. Trace (c) shows the input and output waveforms for square wave response at about 1/3 Hz.
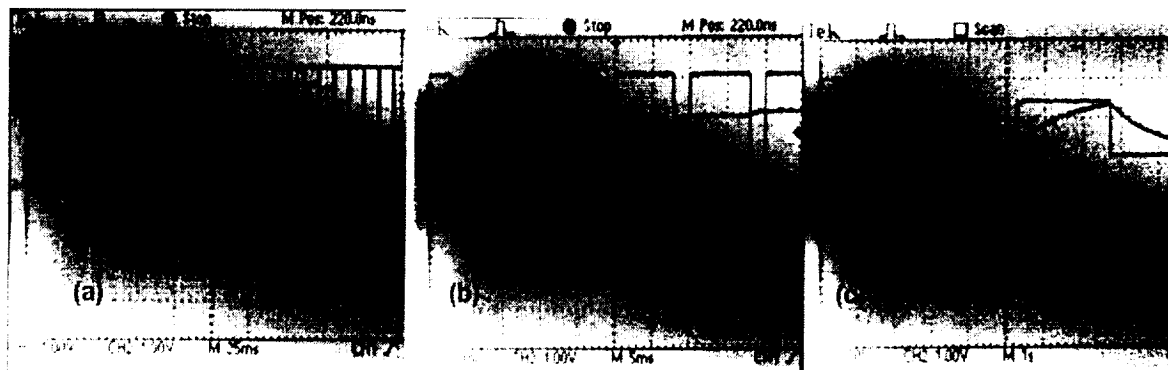

Figure 3: Scope traces for single pole filter, reduced switch rate

4 Component Sharing

If the duty cycle is small, then Kaehler [3] observes active elements, such as an op amp, can be shared. In addition, since the resistor R does not preserve any state information, it can also be shared, either between several independent data channels, or between several stages in a multi-pole filter. This can permit R to be made quite a bit larger than might otherwise be practical, and

lends itself along with the duty cycle ratio improvement to the design of on-chip filters with quite low cutoff frequencies.
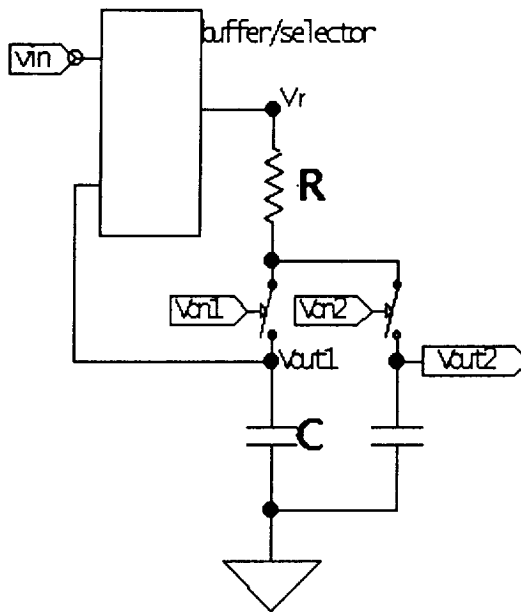

Figure 4: 2-pole filter sharing resistor R

Figure 4 shows a simplified 2-pole filter arrangement sharing a resistor R. Von1 and Von2 are separate non-overlapping "on" duty cycle pulses occurring within the same overall sampling cycle. During Von1, the buffer/selector applies Vin to node Vr, and the output appears at Vout1 as before. When Von1 goes low, Vout1 must remain constant as it is now the input to the next filter stage. During Von2, the buffer/selector applies Vout1 to node Vr, and the final output appears at Vout2. When Von2 goes low, the value at Vout2 is available as a sampled data value for further processing.

Note that the maximum duty cycle in this arrangement is now 50%, and that the second design guideline above is relaxed by that amount.

5 4-Pole Circuit Designs

Going from 2 poles to 4 is a straightforward extension of the above example. Details of the buffer/selector for a single-ended design are shown in Figure 5. An output buffer amp with unity gain is also provided to avoid loading the final capacitor.
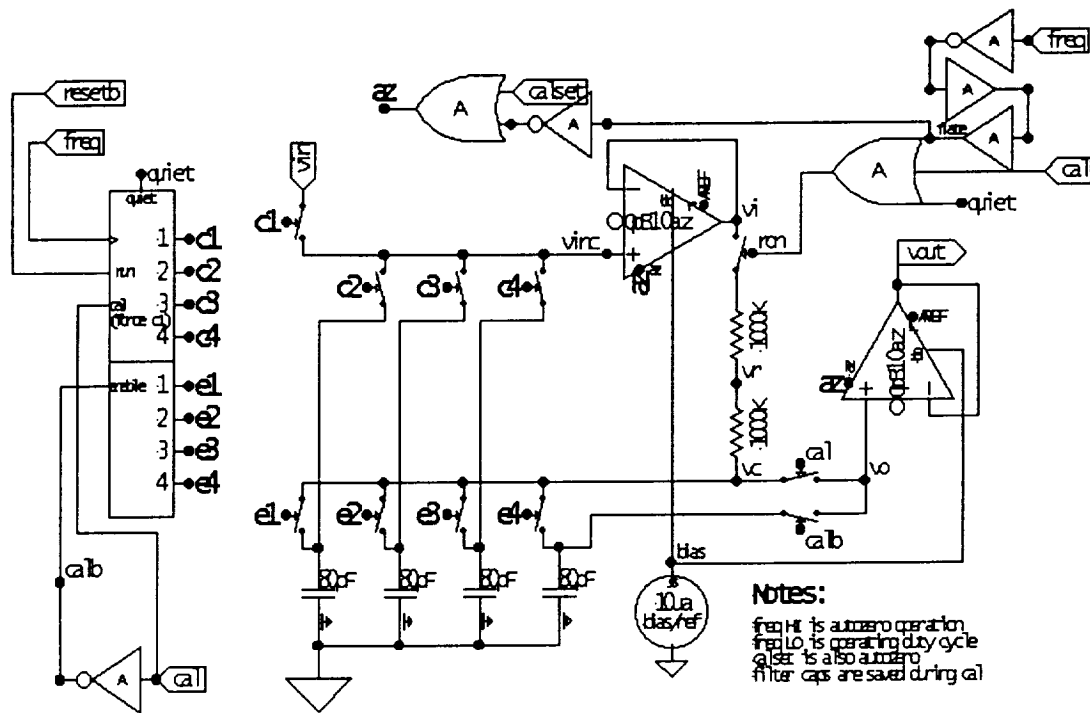
Figure 5: Design of 4-pole filter with auto-zero op amps

Further advantage is taken of the non-continuous operation of the circuit to include an offset auto-zero feature in the op amps controlled by the "cal" signal. To follow normal operation of the circuit, assume "cal" is low and "calb" is high. Switches are "closed" when their control signals are high.

The block on the left contains a simple state machine that generates the separate duty pulses for each filter stage from a single master pulse train.

A fully differential design has also been done, which does not require the auto-zero feature, but requires two resistors and two sets of capacitors. In that case the capacitors were reduced from 80pF to 40pF to conserve layout space. The primary effort in this design was to automatically derive the correct duty cycles from a general fixed 5mHz clock and three bits of filter control information, f1, f2 and fast, rather than directly inputting a continuously variable pulse width. This resulted in a significantly different state machine as shown in Figure 6.
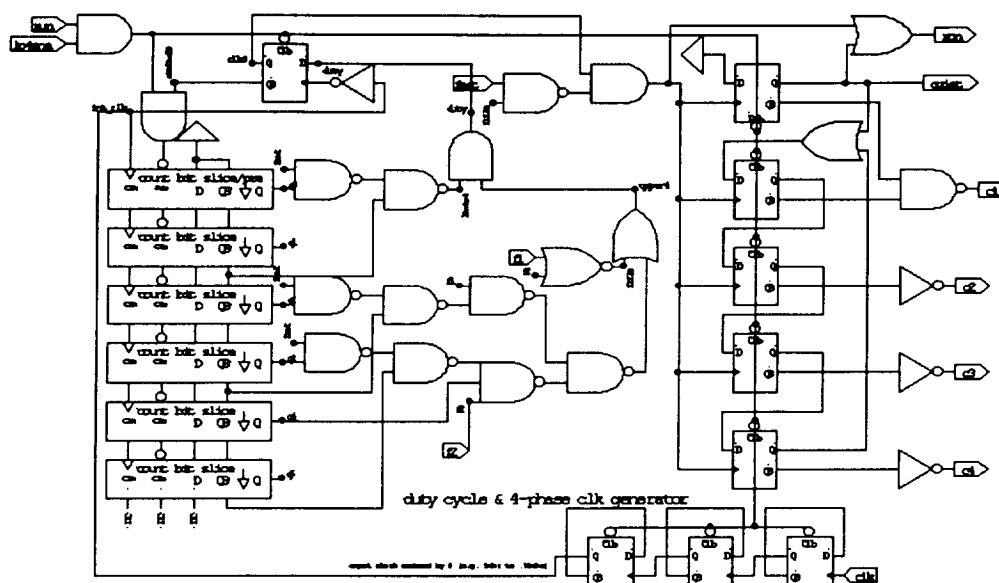
Figure 6: State machine for fully differential 4-pole filter with duty cycle derived from clock
(boxes "count bit slice" contain 1 DFF each with XOR for counting, other boxes are DFF's)

Finally, a fully differential 2-pole design was done for 2.0μ CMOS, with reduced resistor R, and a vastly simplified (single flip flop) state machine.

All designs are produced as auto-routed blocks of standard cells, with guard bars around n-type devices for latchup prevention, and disallowing of cell crossing to avoid coupling switching noise into components such as capacitors where it might result in cumulative error. This is not the most area efficient approach, but allows very rapid prototyping of several designs. A hierarchical block approach is used so that analog and digital circuitry and power/ground can be isolated (which was done extensively in the fully differential design), and the components integrated into system chips without extensive re-verification.

6 Evaluation of Multi-Pole Designs

Presumably due to the use of a new logic family, the 4-pole single ended state machine does not operate as predicted by simulation (there was also poor behavior from an un-related shift register on this chip), and the other two designs have not returned from fabrication, so full evaluation is not available at this time.

7 Op Amp Switching Considerations

As indicated, for small duty cycles op amp settling time could be a source of error. The op amp speed and equalization scheme should be optimized to minimize settling time. This is only a problem in the multi-pole designs, where a buffer op amp is use to dynamically re-configure the circuit. Simulations indicated our amps were adequate to switch between filter stages. After all, full scale step functions must be eliminated from this sampled data circuit by a prior anti-aliasing filter, and signal values in successive stages of the filter are driven closer and closer together.

However, in the auto-zero configuration presented above, the op amp is required to snap into it's auto-zero feedback loop, and snap right back into the signal chain, representing an arbitrary step function of up to 1/2 the full scale range, possibly inducing significant settling time. Any residual unsettled error in the auto-zero mode is incorporated as offset error.

To minimize auto-zero settling time, a second smaller equalization capacitor was provided for each op amp for this mode. Each equalization capacitor retains its state information when the other mode is being used, reducing both the power and time required for switching modes.

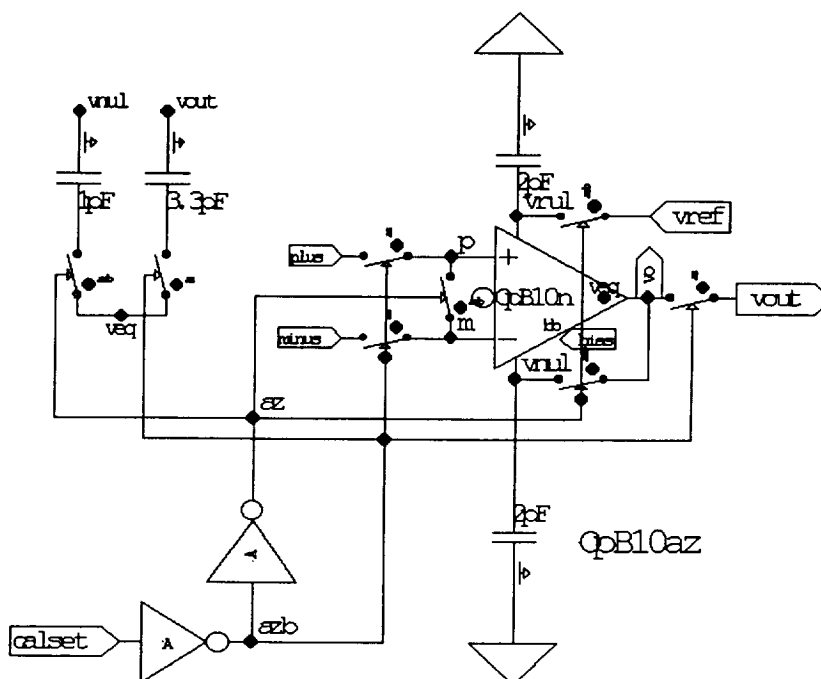Figure 7: Auto-zero op amp with switched EQ

Figure 7 shows an op amp with a second set of input ports, "null ports", used for adjusting offset, and various switches for zeroing the input when in auto-zero mode. The null port adjustment is stored on 2pF capacitors, making it necessary to re-zero several times a second for optimum adjustment. The op amp represented by OpB10n has no internal equalization, but a port "veq" for it to be attached, and in the upper left is shown the separate equalization capacitors and switching arrangement. Simulations indicated that with the high speed logic available in HP 0.5μ, the traditional non-overlapping clocks for this circuit would not be required.

The idea of switched EQ deserves further analysis and better experimental verification, which is planned for a separate investigation to be conducted with collaborators at Mississippi State University.

8 Anti-Alias Considerations

As with any sampled data circuit, a preceeding anti-alias filter is required to maintain signal integrity. The 1 kHz switch rate used in the above experiments might suggest the anti-alias circuit itself would be rather challenging. Indeed it would be, although linear techniques for pushing down cutoff frequency are available [8] and could be used. If in the circuit examples given, a cutoff frequency of only 20 Hz and not 0.2 Hz were the lowest required, the switch rate could be pushed up to 100 kHz, a more tractable configuration for on-chip anti-alias.

Note that in a multi-pole configuration, the effective sample rate is not the switch rate, but 1/N of the switch rate, where N = number of stages in the filter. So for a 4-pole filter configured as shown and switching at 100 kHz, the effective sample rate is 25 kHz.

If the minimum usable "on" time can be reduced, the switch rate can be increase for a given duty cycle. This requires faster settling op amps, possibly extending the switched EQ concept described above to have a separate EQ for each filter stage. This increase in area and complexity will then have to be considered against alternative designs.

9 Conclusions

In general, the performance of the auto-zero op amps has not yet lived up to expectations, which is what motivated the transition to fully-differential designs. In the straightforward transition to full differential, all the circuitry, including R and C's, are duplicated. Simultaneously our desire to operate at higher voltages than 3.3v have induced migration to less dense processes (2.0μ and 1.2μ), with such increase in area for the passive components that the 4-pole design had to be revised downard to 2-poles. Possible fruitful areas for investigation might include some clever topology to share the state capacitors in the differential configuration, and using additional switching topology to share the resistor.

The general impression from working with the circuit is that it is simpler to design and use and probably less area than getting comparable performance from a switched capacitor circuit [7]. However, this needs to be quantitatively investigated. The dependence of cutoff frequency on the absolute values of on-chip passives would be a disadvantage for some applications, but appears acceptable for sensor data acquisition.

While subject to the same sampled data constraints, the circuit is far simpler than working with a digital signal processing (DSP) logic block, and more compact. A DSP would have to provide at least one state variable for each capacitor, and a 10 or 12 bit register in these technologies with wiring is comparable or larger than a 40 or 80pF capacitor. If the DSP algorithm were carefully chosen to avoid multiplying other than by powers of 2, it might be designed in a comparable size. Noise coupling and power dissipation would be interesting issues for comparison with a DSP.

In comparing the technique to continuous time analog (i.e. conventional) circuits which are capable of scaling the cutoff frequency downward such as described in an earlier work [8], those techniques suffer from scaling the op amp offset by the same amount, which is exacerbated further in a multi-pole design. Note that there is an offset degradation of 4x in the switched resistor 4-pole design also, since the signal passes repeatedly through the same op amp,

compared to a degradation of 100x for frequency scaling from something reasonable like 2000 Hz down to a modest 20 Hz for the continuous time filter. And the switched resistor filter is capable of much greater scaling.

The early literature on switched resistors was of a general nature, and was overtaken by switched capacitor technology. Either technology needs to be optimized for specific applications before becoming useful. With increasing CMOS switching speeds (CMOS was not even available before switched resistor essentially died), and a requirement to provide simple designs having very low frequency operation integrated within sensor data acquisition chips, we believe switched resistor may finally find some practical application. The ability to share not only the active components, which have become vanishingly small, but also some of the large passive components as documented in this paper, provides a further advantage not previously available.

References

[1] H. C. Lin, W. R. Davis, and Y. Sun, "Active and passive multiplication of RC time constants for subaudio frequency integrated filters," 1966 IEEE Region 6th Ann. Conf. Rec., vol 1, pp.1-9.

[2] Y. Sun and I. T. Frisch, "Resistance Multiplication in Integrated Circuits by Means of Switching," IEEE Transactions on Circuit Theory, vol. CT-15, pp. 184-192, September 1968.

[3] J. A. Kaehler, "Periodic-Switched Filter Networks - A Means of Amplifying and Varying Transfer Functions," IEEE Journal of Solid-state Circuits, vol. SC-4, pp. 225-230, August 1969.

[4] L. T. Bruton and R. T. Pederson, "Tunable RC- Active Filters Using Periodically Switched Conductances," IEEE Transactions on Circuit Theory, vol. CT-20, pp. 294-301, May 1973.

[5] P. M. Embree, "The Analysis and Performance of a Tunable State Variable Filter Employing Switched Resistor Elements," Master's Thesis in Electrical Engineering, University of Illinois at Urbana-Champaign, 1982.

[6] R. L. Geiger, P. E. Allen and D. T. Ngo, "Switched-Resistor Filters - A Continuous Time Approach to Monolithic MOS Filter Design," IEEE Transactions on Circuits and Systems, vol. CAS-29, pp. 306-315, May 1982.

[7] Jan-Michael Stevenson, Gunhee Han, and Edgar Sanchez-Sinencio, "A Low-Noise, Programmable Gain and Cutoff Frequency Data Acquisition Chip: Version Two", Texas A&M Univ., 1996, Dept. of EE, College Station, TX 77843-3128.

[8] R. Shuler, "Single-Chip Instrumentation Module for Radiation Environment", 7th Annual NASA/IEEE VLSI Symposium, October 1998, University of New Mexico, Albuquerque, NM.